On Making Encrypted Web Traffic Resistant to Timing-Analysis Attacks

Saman Feghhi and Douglas J. Leith
Trinity College Dublin, Ireland

arXiv:1610.07141v1 [cs.NI] 23 Oct 2016

***Abstract*—In this paper we introduce a trace-based tunnel that is resistant to traffic analysis in the sense that it provides deniability to users that a specific web page was fetched given that a packet trace is observed on the tunnel. We present a scheduler design for managing the transmission of traces to satisfy user traffic demand while maintaining reasonably low delay and throughput overhead due to dummy packets. Experimental results are also presented demonstrating the effectiveness of this scheduler under a range of realistic network conditions and actual web page fetches.**

***Index Terms*—Timing-analysis attacks, Traffic analysis defence, Website fingerprinting, Network privacy, VPN.**

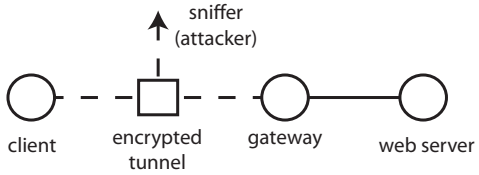

Fig. 1: Schematic illustrating attacker of the type considered. A client machine is connected to an external network via an encrypted tunnel (ssh, SSL, IPSec etc.). The attacker can sniff packets traversing the tunnel, but has no other information about the clients activity.

I. Introduction

Encrypting traffic to protect it from eavesdroppers is one of the basic building blocks of secure networks. However, while encryption conceals the contents of transmitted packets other features of the transmitted packet stream, such as the packet sizes and timings, often still contain revealing information. Taking advantage of this, over the last 10 years a number of increasingly powerful attacks have been demonstrated against encrypted packet streams, and encrypted web traffic in particular. While attacks against HTTPS were initially based on packet sizes and counts, attacks using only packet timing information have recently been demonstrated, *e.g.* [8]. The importance of the latter is that they are immune to defences currently being rolled out, such as padding packets to make them all the same size, plus they do not require *a priori* knowledge of the start and end of each web fetch. Such developments strongly motivate revisiting how we transmit encrypted traffic.

We consider an attacker of the type illustrated in Figure 1. The attacker can sniff packets traversing the encrypted tunnel[1] and the attacker's objective is to use this information to guess, with high probability of success, the web pages which the client visits. We are not concerned with concealing the fact that the user is browsing the web but rather with concealing the particular web pages visited. Therefore, given an observed packet sequence on the network we would like there to be a sufficiently large number of different web fetches, or combinations of web fetches, that could generate the same packet sequence with reasonable probability. When this holds then a user can reasonably deny that they have fetched any particular page, claiming that instead they fetched a different page or combination of pages. Privacy is therefore embodied in the *indistinguishability* of sequences of web fetches given with regard to the packet sequences which they generate. There is also a need to ensure diversity, in an appropriate sense, amongst the sequences of web fetches that can explain an observed packet sequence. For example, if only combinations of fetches of pages from the same web site could generate an observed packet sequence then privacy is undermined. However, given the huge number of web pages in the Internet lack of diversity seems unlikely to be a major concern.

Indistinguishability directly targets one of the weaknesses of existing attacks. Namely, the assumption that web fetches occur sequentially, one at a time. When this assumption is violated then the attacker needs to estimate the mixture of web pages fetched. By ensuring that sufficiently many mixtures of fetches generate the same packet trace then we can ensure that this inverse estimation task is ill-posed unless the attacker has strong *a priori* knowledge of the web pages fetched (but such a powerful attacker has little to gain by inspecting packet traces).

The main contributions of the paper are as follows: (i) we introduce a trace-based tunnel that is resistant to traffic analysis in the sense that it provides deniability to users that a specific web page was fetched given that a packet trace is observed on the tunnel, (ii) we present a scheduler design for managing the transmission of traces to satisfy user traffic demand while maintaining reasonably low delay and throughput overhead due to dummy packets, and (iii) we build an experimental prototype of the traffic shaping tunnel and scheduler and carry out an extensive performamnce evaluation that demonstrates its effectiveness under a range of network conditions and real web page fetches.

Work supported by SFI grant 11/PI/1177 and 13/RC/2077.

[1]We focus on tunnels for concreteness and due to the importance of VPNs, but almost all of our ideas and results carry over to HTTPS with little change.

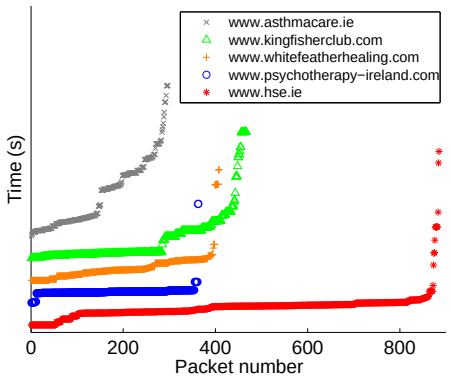

Fig. 2: Time traces of uplink traffic from 5 different Irish health-related web sites are shown. It can be seen that the web site time traces exhibit distinct patterns. The traces are shifted vertically to avoid overlap and facilitate comparison. Figure reprinted from [8].

II. Preliminaries

A. Anatomy of a Web Page Fetch

Before considering defences against timing analysis attacks we briefly review the reasons why packet timing information can be highly informative. When traffic is encrypted the packet source and destination addresses and ports and the packet payload are hidden and the packets may be padded to be of equal size, so that packet size information is also concealed. An attacker sniffing such encrypted traffic is therefore able only to observe the direction and timing of packets. Figure 2 plots the timestamps of the uplink packets sent during the course of fetching five different health-related web pages (see below for details of the measurement setup). The x-axis indicates the packet number k within the stream and the y-axis the corresponding timestamp. It can be seen that these timestamp traces are distinctly different for each web site.

To gain insight into the differences between the packet timestamp sequences in Figure 2, it is helpful to consider the process of fetching a web page in more detail. To fetch a web page the client browser starts by opening a TCP connection with the server indicated by the URL and issues an HTTP GET or POST request to which the server then replies. As the client parses the server response it issues additional GET/POST requests to fetch embedded objects (images, css, scripts *etc.*). These additional requests may be to different servers from the original request (*e.g.* when the object to be fetched is an advert or is hosted in a separate content-delivery network), in which case the client opens a TCP connection to each new server in order to issue the requests. Fetching of these objects may in turn trigger the fetching of further objects. Note that asynchronous fetching of dynamic content using, *e.g.* AJAX, can lead to a complex sequence of server requests and responses even after the page has been rendered by the browser. Also, typically the TCP connections to the various servers are held open until the page is fully loaded so that they can be reused for later requests (request pipelining in this way is almost universally used by modern browsers). It is this sequence of events while fetching the objects in a page that creates a packet timing "signature" which can be used to deanonymise the encrypted traffic and estimate the web page fetched with high probability. For example, in [8] web pages are correctly identified with $> 90\%$ success rate.

B. Related Work

The general topic of traffic analysis has been the subject of much interest, and a large body of literature exists. Some of the earliest work specifically focussed on attacks and defences for encrypted web traffic appears to be that of Hintz [12], which considers the SafeWeb encrypting proxy. In this setup (i) web page fetches occur sequentially with the start and end of each web page fetch known, and for each packet (ii) the client-side port number, (iii) the direction (incoming/outgoing) and (iv) the size is observed. A web page signature is constructed consisting of the aggregate bytes received on each port (calculated by summing packet sizes), effectively corresponding to the number and size of each object within the web page.

Subsequently, Bissias *et al* [1] considered an encrypted tunnel setup where (i) web page fetches occur sequentially with the start and end of each web page fetch known, and for each packet (ii) the size, (iii) the direction (incoming/outgoing) and (iv) the time (and so also the packet ordering) is observed. The sequence of packet inter-arrival times and packet sizes from a web page fetch is used to create a profile for each web page in a target set and the cross correlation between an observed traffic sequence and the stored profiles is then used as a measure of similarity.

Most later work has adopted essentially the same model as [1], making use of packet direction and size information and assuming that the packet stream has already been partitioned into individual web page fetches. For example in [21] the timing information is not considered in the feature set, hence the attack can be countered with defences such as BuFLO in [6] leading to a success rate of only 10%. In [13], [11] Bayes classifiers based on the direction and size of packets are considered while in [18] an SVM classifier is proposed. In [14] classification based on direction and size of packets is studied using Levenshtein distance as the similarity metric, in [16] using a Gaussian Bag-of-Words approach and in [21] using K-NN classification. In [2] using a SVM approach a classification accuracy of over 80% is reported for both SSH and Tor traffic and the defences considered were generally found to be ineffective. Similarly, [6] considers Bayes and SVM classifiers and finds that a range of proposed defences are ineffective. In [10] remote inference of packet sizes from queueing delay is studied.

Recently, powerful attacks based on using packet timing information alone have been demonstrated [8]. Such attacks are immune to defences based on packet padding etc.

Among the earliest defences against traffic analysis attacks is HTTPOS [15] and the work of Wright et al [24]. These modify packet size, packet timing and payload size and provide fairly successful defence against a number of classifiers [20], [1], [11], [25], [5]. Similarly in [6], BuFLO is proposed as a defence against traffic analysis attacks. The

model combines previous countermeasures such as use of fixed packet sizes and constant rate traffic to respond to the attacks based on packet size information. Subsequently, CS_BuFLO in [3], made modifications to BuFLO based on rate adaptation, but this comes at the expense of a high bandwidth overhead. Similarly, the Tamaraw defence in [4] works similarly to the BuFLO concept, with the length of packets and length of traces as design parameters.

Similarly to the present work, the idea of sending super-traces that represent a group of web pages is considered in [17]. However the features considered are based on packet size and frequency whereas here we consider the timing features of a packet sequence. More recently Wang et al [23] suggested a half-duplex transmission method where the client cannot send multiple requests to servers in parallel.

III. Achieving Privacy

A. Indistinguishability

As already noted, we are not concerned with concealing the fact that the user is browsing the web but rather with concealing the particular pages visited. Privacy is therefore embodied in the indistinguishability of sequences of web fetches given with regard to the packet sequences which they generate.

It is important to note that we do not insist that fetches of single web pages individually be indistinguishable. Rather, we require that the observed packet sequence corresponding to the fetch of each web page can be reasonably explained by other combinations of web fetches. So, for example, the packet sequence generated by the fetch of a large web page could equally have been generated by sufficiently many sequences of fetches of small web pages. Fetches of single web pages must still be indistinguishable from one another when they cannot plausibly be generated by combinations of other page fetches. That is, we require fetches of sufficiently small web pages to be atomic in the sense that they are indistinguishable from one another.

B. Defence By Use of "Traces"

The basic problem with the packet sequence generated by a web fetch is that it contains many packets (often several hundred, frequently more) and so amounts to an observation of a point in high dimensional space. It is increasingly recognised that high dimensional data carries a major risk of de-anonymisation since data points tend to be sparsely distributed in high dimensions (each point can be individually distinguished). Our approach is therefore to reduce the dimension of the observed packet sequence data. We do this by gathering packets into larger groups and transmitting these groups in a uniform way.

This approach is illustrated in Figure 3a. Here, a web fetch starts at time t_0. Dummy packets (indicated in grey) are inserted so that the observed sequence of packets, which is a combination of dummy plus user packets, has a uniform profile and duration. We refer to this uniform profile as a "trace". The rate and duration of the trace are design parameters, that we will return to later[2]. In this example the web fetch completes by time $t_0 + T$ and so fits within a single trace. However, if the duration exceeded T then a second trace would be started so as to mask the additional user packets, see upper schematic in Figure 3b. Similarly, if the rate at which user packets arrive much exceeds the rate of a trace then two traces can be started in parallel, see lower schematic in Figure 3b. With this approach the observed packet sequence now consists of a sequence of traces.

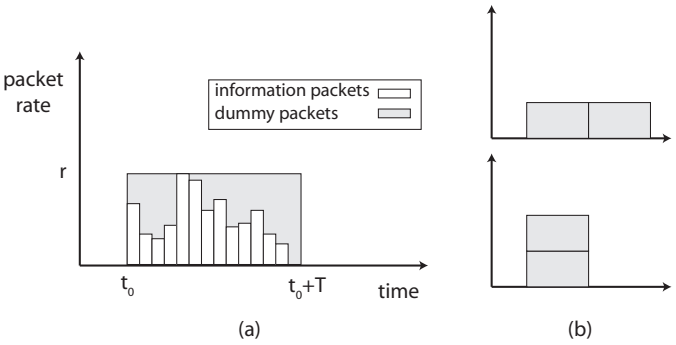

Fig. 3: Illustrating use of dummy packets to force the number and timing of transmitted packets to conform to an uninformative pattern or "trace". When insufficient information packets are available, dummy packets are inserted as needed. When too many information packets are available they are buffered until a transmission opportunity becomes available.

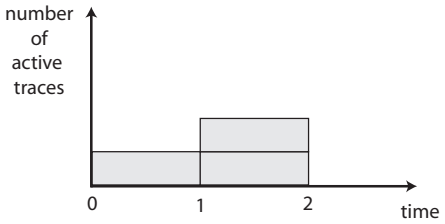

Fig. 4: Example illustrating an observed sequence of traces.

Given an observed sequence of such traces, the indistinguishability question becomes: how many different sequences of web fetches can reasonably generate the observed trace sequence. For example, consider the sequence of traces shown in Figure 4. Time is slotted, with each slot corresponding to the duration of a trace. Traces start at the beginning of a slot, and suppose a maximum of n traces are possible in each slot (e.g. limited by the link capacity). Suppose we have a set W of web pages of interest. Of these, fetches for web pages $W_1 \subset W$ can be covered by a one trace, $W_1 \subset W$ by two traces and so on. Let p_w be the probability that web page $w \in W$ is fetched and assume, for simplicity, that pages are fetched independently. Let X_w be a random variable which takes value 1 when web page w is fetched and 0 otherwise. Let $T \subset \{0, n\}^k$ denote the observed sequence of traces of

[2]We might also use a trace where the rate is not constant over time and also allow the trace used to be drawn from a pre-defined set of traces.

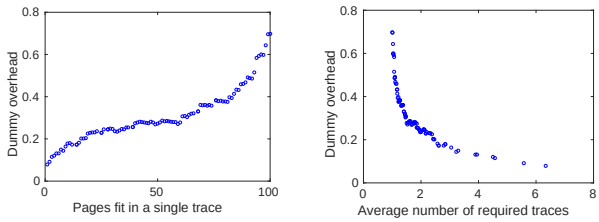

(a) (b)

Fig. 5: Illustrating the overhead of dummy packets vs trace duration. Data is for Alexa top 100 web pages.

length k slots, with $T = (1, 2, 0, \ldots, 0)$ in Figure 4 . Then,

$$\mathrm{P}(X_w = 1|T, w \in W_1) = 1 - \prod_{j=1}^{k} \left(1 - \frac{p_w}{\sum_{i \in W_1} p_i}\right)^{T_j} \quad (1)$$

and similarly we can calculate $\mathrm{P}(X_w = 1|T, w \in W_2)$ etc. In the example in Figure 4 suppose $p_w = 1/|W|$. Then for a web page $w \in W_1$ that fits inside a single trace the probability that it was fetched is $1 - \left(1 - \frac{1}{|W_1|}\right)\left(1 - \frac{1}{|W_1|}\right)^2$. For example, when $|W_1| = 100$ this probability is 0.03 and when $|W_1| = 1000$ the probability is 0.003 and it can be seen that provided the cardinality of W_1 is reasonably large a user can with high plausibility deny that they fetched a given web page.

Note that in the above example an attacker can infer that the user did not fetch any web pages that cannot fit inside three traces or fewer. However, given the extremely large number of possible web pages and the fact that most of these pages are not fetched by a given user we argue that the absence of a fetch for a web site is much less informative to an attacker than the presence of a fetch.

We note also that the option exists to inject dummy traces into the tunnel link to add a further level of deniability. Using such an approach the potential exists to formulate indistinguishability as a form of differential privacy. Namely, such that the addition of the fetch of any individual web page has limited impact on the sequence of traces observed on the link. However, we leave this as future work and do not pursue it further here.

C. Rate and Duration of a Trace

By buffering user packets at the tunnel ingress until enough are available to fill a trace we can avoid the overhead of sending dummy packets. However, this comes at the cost of increased delay, perhaps much increased delay. However, in modern networks quality of service is frequently achieved by over-provisioning of network capacity and as a result bandwidth is often relatively plentiful. Conversely, users are known to be sensitive to delay when using online services. For example, Amazon estimates that a 100ms increase in delay reduces its revenue by 1% [9], Google measured a 0.74% drop in web searches when delay was artificially increased by 400ms [26] while Bing saw a 1.2% reduction in per-user revenue when the service delay was increased by 500ms [19]. Hence, in order to enhance resistance to traffic analysis attacks it is preferable to sacrifice some bandwidth by sending dummy packets rather than incurring excessive delay.

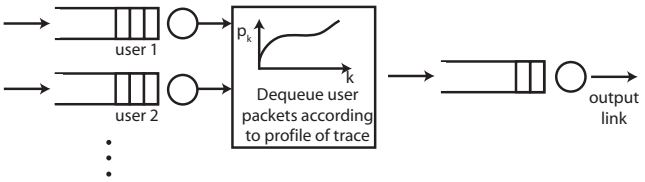

Fig. 6: Schematic illustrating the VPN gateway setup considered.

To gain some insight into the overhead of dummy packets associated with different durations of trace we fetched the home pages from the Alexa top 100 finance and health web sites in Ireland. Figure 5a plots the average fraction of packets in a trace which are dummy packets vs the number of web sites that fit inside the trace. A trace of duration 10890ms is needed to cover all 100 web sites and a trace of duration 4947ms to cover 50 web sites. As might be expected, the fraction of dummy packets increases as the duration of the trace is increased to include more web pages. When all 100 web pages fit inside a single trace the overhead is around 70% but this falls to around 30% for a trace that covers 50 web pages. Figure 5b plots the fraction of dummy packets vs the maximum number of consecutive traces used to cover all 100 web pages (so the value for one trace corresponds to the data in Figure 5a). It can be seen that as the number of traces used increases the overhead falls. However, use of smaller traces reduces the level of indistinguisability provided and so a trade-off exists between privacy and dummy packet overhead.

Since it is hard to analytically quantify the trade-off between privacy, dummy packet overhead and delay (the delay aspect is especially difficult to analyse mathematically), we will revisit this trade-off shortly using experimental data.

IV. Scheduling Traces

A privacy-enhanced tunnel using the trace approach must schedule the start-up of traces so as to minimise user delay and dummy packet overhead. The design and analysis of an efficient scheduler for this is the main technical challenge that needs to be solved in order to construct such a tunnel and is addressed in this section.

A. Network Setup

The setup considered is illustrated in Fig 6. Let $\mathcal{U} = \{1, 2, \cdots, n_u\}$ denote the set of users, $a_k^{(u)} \in \mathbb{N}$ denote the number of packet arrivals for user u at time k and $\bar{a}^{(u)} := \lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} a_k^{(u)}$ the average number of packet arrivals. Packets for each user u are held in separate queues, with the queue occupancy for user u at time k being denoted by $q_k^{(u)}$. Letting $w_k^{(u)} \in \mathbb{N}$ be the number of packets dequeued at time k then $q_{k+1}^{(u)} = [q_k^{(u)} + a_k^{(u)} - w_k^{(u)}]^+$ where the notation $[\cdot]^+$ has the usual meaning.

User packets are dequeued according to a predefined *trace*. Multiple traces may be active simultaneously in order to

increase the rate at which user packets are dequeued. A trace f is a sequence $p_j^{(f)} \in \mathbb{N}$, $j = 1, 2, \ldots, n^{(f)}$ where $p_j^{(f)}$ is the number of packets to be transmitted at time j and $n^{(f)}$ is the duration of the trace. Importantly, when no user packets are available to send at time j, then dummy packets will be transmitted to ensure that $p_j^{(f)}$ packets are always transmitted. Let $P^{(f)} = \sum_{j=1}^{n^{(f)}} p_j^{(f)}$ denote the total number of packets in trace f. In general we might have a family of m different traces $\mathcal{F}$ to choose from.

Packets from a trace are queued at the output link before transmission over the tunnel, see Fig 6. As is usual in Internet links we leave rate control to the end hosts. If the aggregate send rate persistently exceeds the output link capacity then the queue at this link will eventually overflow and cause packet loss which end hosts can then use as a congestion indicator, e.g. by using TCP congestion control.

New traces from family $\mathcal{F}$ are started as needed to service user packet arrivals, and multiple copies of a trace may be active at the same time so as to increase the sending rate if needed. Indexing these active traces by $1, 2, \cdots$ let $\mathcal{T}$ denote the set of trace indices, τ_t the start time of trace $t \in \mathcal{T}$ and f_t the member of family $\mathcal{F}$ to which the trace corresponds. At time k the set of active traces is $\mathcal{T}_k := \{t \in \mathcal{T} : k \in \{\tau_t, \ldots, \tau_t + n^{(t)}\}\}$ and the number of packets that are transmitted at time k is $\sum_{t\in\mathcal{T}_k} p_{k-\tau_t}^{(f_t)}$, therefore $\sum_{u\in\mathcal{U}} w_k^{(u)} \leq \sum_{t\in\mathcal{T}_k} p_{k-\tau_t}$. The average rate of packet transmissions can be no more than the capacity c in packets/slot of the outgoing link i.e. $\lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \sum_{t\in\mathcal{T}_k} p_{k-\tau_t}^{(f_t)} < c$.

Since $\sum_{t\in\mathcal{T}_k} p_{k-\tau_t}^{(f_t)}$ packets are sent then if there are insufficient user packets available to send at time k we need to send $\sum_{t\in\mathcal{T}_k} p_{k-\tau_t}^{(f_t)} - \sum_{u\in\mathcal{U}} w_k^{(u)}$ dummy packets. Our task is to activate traces and thereby adjust the number of user packets transmitted $w_k^{(u)}$ so as to minimise the number of dummy packets sent while still servicing all of the user packets in a timely manner.

To simplify the presentation that follows we will confine ourselves to setup with a single type of trace, i.e. family $\mathcal{F}$ contains a single member, the extension to multiple types of trace being straightforward. We take advantage of this to streamline notation, shortening $p_k^{(f_t)}$ to p_k and $n^{(t)}$ to n.

B. Synchronised Scheduler

To proceed we first construct an approximation to the above problem setup where active traces start/stop in a synchronised fashion. Traces started in slot 1 finish at slot n (since traces are of n slots duration), traces started at slot $n+1$ finish at slot $2n$ and so on. Hence, by restricting ourselves to starting traces at times $1, n+1, 2n+1, \ldots$ we can ensure synchronised start/stop of traces and avoid partial overlapping of traces. We will relax this restriction later, but the absence of overlapping traces simplifies the initial analysis and assists with gaining insight into the design issues to be considered.

Formally, partition time slots $k = 1, 2, \ldots$ into groups of n slots, see Fig 7. Denote the first group of slots by $\mathcal{G}_1 := \{1, \ldots, n\}$, the second by $\mathcal{G}_2 := \{n+1, \ldots, 2n\}$ and so on. Suppose a new trace can only be activated at the start of a group of slots. Since each group is n slots long each trace will therefore have finished by the end of the group of slots in which it started. Letting y_g denote the number of traces active in group g, since the traces within a group are all activated at the same time the number of packets that are transmitted in slot $k \in \mathcal{G}_g$ simplifies to $y_g p_{k-(g-1)n}$.

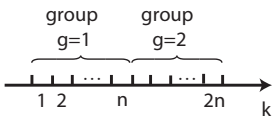

Fig. 7: Gathering slots into groups of n consecutive slots. In the synchronised case a trace can only be activated in the first slot of a group, and will have completed by the last slot in that same group since traces are of duration n slots.

Let $b_g^{(u)} := \sum_{k\in\mathcal{G}_g} a_k^{(u)}$ be the number of packet arrivals for user u in group of slots g and $\bar{b}^{(u)} := \lim_{K\to\infty} \frac{1}{K} \sum_{g=1}^{K} b_g^{(u)} = \bar{a}^{(u)} n$ be the average number of packet arrivals for user u. Then provided the average number of packets served $\lim_{K\to\infty} \frac{1}{K} \sum_{g=1}^{K} \sum_{k\in\mathcal{G}_g} w_k^{(u)} = \bar{y}P > \sum_{u\in\mathcal{U}} \bar{b}^{(u)}$ then the arriving user packets can all be served.

1) Convex Optimisation: To minimise the number of dummy packets sent while serving the user packet arrivals we end up with the following formal convex optimisation P,

$$\min_{0\leq\bar{y}\leq y_{max}} \gamma\bar{y}$$
$$s.t. \quad \sum_{u\in\mathcal{U}} \bar{b}^{(u)} \leq \bar{y}P$$

where $\gamma \geq 0$. This optimisation is formal in the sense that we can see that the solution is trivially $\bar{y}^* = \frac{1}{P} \sum_{u\in\mathcal{U}} \bar{b}^{(u)}$ provided $\frac{1}{P} \sum_{u\in\mathcal{U}} \bar{b}^{(u)} \leq y_{max}$. However, we would like to derive an algorithm that finds the solution without requiring a priori knowledge of the mean user packet arrival rates $\bar{b}^{(u)}$ and we would also like to ensure the stability of the queue in the scheduler, and it is for these reasons that the formulation as an optimisation will prove useful.

2) Online Solution: The Lagrangian for the above optimisation is $L(\bar{y}, \lambda) := \gamma\bar{y} + \lambda\left(\sum_{u\in\mathcal{U}} \bar{b}^{(u)} - \bar{y}P\right)$. By the KKT conditions it can be seen that the multipliers satisfies $\lambda - \gamma/P = 0$ at an optimum. Suppose the Slater condition is satisfied and so strong duality holds. Then a solution can be found using the following dual subgradient update

$$x_g \in \arg\min_{x\in[0,y_{max}]} \partial L(\bar{y}_g, \lambda_g) x \tag{2}$$
$$= \arg\min_{x\in[0,y_{max}]} (\gamma - \lambda_g P) x \tag{3}$$
$$\bar{y}_{g+1} = (1-\beta)\bar{y}_g + \beta x_g \tag{4}$$
$$\lambda_{g+1} = [\lambda_g + \alpha(\sum_{u\in\mathcal{U}} \bar{b}^{(u)} - \bar{y}_g P)]^+ \tag{5}$$

with $\alpha > 0$, $\lambda_1 = 0$ and $0 \leq \beta < 1$. Update (3)-(4) is a Frank-Wolfe descent update while (5) is a dual subgradient ascent update and the joint convergence of these to a ball around the optimum is established by, for example, Corollary 2 in [22].

However, since the Lagrangian is linear in x_g the solution x_g in (3) is either 0 or y_{max} i.e. we end up with a bang-bang

type of solution whose running average $\bar{y}_g$ converges to the optimum. Such bang-bang scheduling of packets is not well suited to Internet applications since it is bursty and can lead to large fluctuations in delay. We therefore consider the following modifcation to the update:

$$\delta\bar{y}_g \in \arg\min_{x\in[-1,1]}(\gamma-\lambda_g P)x \tag{6}$$

$$\bar{y}_g = [\bar{y}_{g-1}+\beta\delta\bar{y}_g]^{[0,y_{max}]} \tag{7}$$

$$\lambda_{g+1} = [\lambda_g+\alpha(\sum_{u\in\mathcal{U}}\bar{b}^{(u)}-\bar{y}_g P)]^+ \tag{8}$$

where $[y]^{[0,y_{max}]}=y$ when $y\in[0,y_{max}]$, 0 when $y<0$ and y_{max} when $y>y_{max}$.

Observe that $|\lambda_{g+1}-\lambda_g|\le\alpha|\sum_{u\in\mathcal{U}}\bar{b}^{(u)}-\bar{y}_g P|\le\alpha b^\diamond$ where $b^\diamond:=\sum_{u\in\mathcal{U}}\bar{b}^{(u)}$. The following lemma now establishes that update (6)-(7) ensures that λ_g is bounded:

Lemma 1 (Boundedness)**.** *Consider update (6)-(7). Suppose* $|\lambda_{g+1}-\lambda_g|\le\epsilon/(2(Py_{max}-b^\diamond))$, $\epsilon>0$, *and* $y_{max}>b^\diamond$ *is sufficiently large. Then as* $g\to\infty$ *we have* $|\gamma-\lambda_g P|\le\epsilon':=(2\beta+1)\epsilon/\beta$.

Proof: See Appendix. ∎

Selecting step size $\alpha\le\epsilon/(2Py_{max}(2Py_{max}-b^\diamond))$ is sufficient to satisfy the conditions of this lemma[3]. We have the following immediate consequence. Observe from (8) that,

$$\lambda_{g+1}\ge\lambda_g+\alpha(\sum_{u\in\mathcal{U}}\bar{b}^{(u)}-\bar{y}_g P) \tag{9}$$

and so when $\lambda_1=0$ it follows that

$$\frac{1}{g}\sum_{i=1}^{g-1}(\sum_{u\in\mathcal{U}}\bar{b}^{(u)}-\bar{y}_g P)=\sum_{u\in\mathcal{U}}\bar{b}^{(u)}-y_g^\diamond P\le\frac{\lambda_g}{\alpha g} \tag{10}$$

where $y_g^\diamond:=\frac{1}{g}\sum_{i=1}^{g-1}\bar{y}_g$. Since, by Lemma 1, λ_g is bounded then $\frac{\lambda_g}{\alpha g}\to 0$ as $g\to\infty$ and therefore $\sum_{u\in\mathcal{U}}\bar{b}^{(u)}\le y_g^\diamond P$ as $g\to\infty$. That is, the mean service rate provided by the traces is sufficient to serve the mean rate of user packet arrivals.

Update (6)-(8) requires knowledge of the long-term average arrival rate $\bar{b}^{(u)}$ to carry out step (8), which we do not have. Also, the average number of active traces $\bar{y}_g$ is real valued whereas the number of active traces, which is the quantity needed in order to implement a scheduler, is integer valued. To address these issues we therefore consider the following update:

$$\delta y_g \in \arg\min_{x\in[-1,1]}(\gamma-\tilde{\lambda}_g P)x \tag{11}$$

$$y_{g+1} = [y_g+\delta y_g]^{[0,y_{max}]} \tag{12}$$

$$\tilde{\lambda}_{g+1} = [\tilde{\lambda}_g+\alpha(\sum_{u\in\mathcal{U}}b_g^{(u)}-y_g P)]^+ \tag{13}$$

Observe that the packet arrivals $b_g^{(u)}$ of user u in interval g are now used in (13) and so knowledge of the mean arrival rate $\bar{b}^{(u)}$ is *not* required. Also, the number of active flows y_g is now used rather than the average $\bar{y}_g$ (although observe that this just amounts to selecting $\beta=1$). It follows immediately from Lemma 1 that $\tilde{\lambda}_g$ is bounded provided the packet arrivals $\sum_{u\in\mathcal{U}}b_g^{(u)}$ in an interval are bounded, y_{max} is sufficiently large and step size α is selected sufficiently small.

3) Discussion: The algorithm (11)-(13) is intuitive. Multiplier $\tilde{\lambda}_g$ measures the accumulated mismatch between the arrival of user packets and service of packets ($\sum_{u\in\mathcal{U}}b_g^{(u)}$ is the number of packet arrivals in interval g, y_g is the number of traces active in interval g and $y_g P$ is therefore the number of packets served over the interval). When this mismatch grows too large and exceeds γ/P then we increase the number of active traces, when the mismatch falls to less than γ/P we decrease the number of active traces. That is, the scheduler uses a threshold rule.

As we will shortly see, the value of $\tilde{\lambda}_g$ is closely related to the occupancy of the queues at the input to the scheduler (at the left-hand side of Fig 6) and so in effect the scheduler operates by activating new traces when the queue backlog becomes too large. This also suggests that there is a trade-off between delay and optimality (i.e. throughput efficiency). That is, a large queue backlog means that there are plenty of user packets available to be scheduled and hardly any dummy packets are needed, whereas a smaller queue means the scheduler is more likely to need to insert dummy packets in order to ensure transmission obey the specified trace profile. However, a large queue implies large delay.

4) Queue Stability: Update (11)-(13) stabilises the input queue to the scheduler. The occupancy of the input queue to the scheduler obeys the following update,

$$q_{gn+1} = [q_{gn}+\sum_{u\in\mathcal{U}}b_g^{(u)}-y_g P]^+ \tag{14}$$

at the time slots gn, $g=1,2,\ldots$ corresponding to the end of each group g of slots (so we are looking the queue occupancy in embedded time). Observe that letting $\tilde{q}_g:=\tilde{\lambda}_g/\alpha$ then by update (13),

$$\tilde{q}_{g+1} = [\tilde{q}_g+\sum_{u\in\mathcal{U}}b_g^{(u)}-y_g P]^{[0,y_{max}]} \tag{15}$$

which, apart from time dilation by n and the upper limit y_{max}, is identical to (14). Hence, when y_{max} is sufficiently large we can identify $\tilde{\lambda}_g/\alpha$ with q_{gn} and, by Lemma 1,

$$q_g\le\frac{\gamma+\epsilon'}{P\alpha} \tag{16}$$

for g sufficiently large. Hence, the input queue is uniformly bounded (and so stable). Further, it can be seen that the bound on the queue occupancy is proportional to γ and $1/\alpha$, so decreasing threshold parameter γ or increasing step size α decreases the queue occupancy, but inversely proportional to the trace length P, so longer traces reduce the queue occupancy.

C. Unsynchronised Scheduler

We now relax the synchronisation assumption made in the previous section and allow new traces to start in any slot. This means that traces can, for example, now partially overlap. The payoff is that the potential exists to exploit this extra freedom to achieve improved performance, especially improved delay

[3]Since $|\lambda_{g+1}-\lambda_g|\le\alpha|\sum_{u\in\mathcal{U}}\bar{b}^{(u)}-\bar{y}_g P|\le\alpha\max\{b^\diamond,y_{max}P\}\le\epsilon y_{max}P/(2Py_{max}(Py_{max}-b^\diamond))\le\epsilon/(2(Py_{max}-b^\diamond))$.

performance as a result of being able to start new traces in a more timely way.

As our baseline unsynchronised scheduler we use the following update,

$$\delta x_k \in \arg\min_{x\in[-1,1]} (\gamma - \hat{\lambda}_k P)x \tag{17}$$

$$x_{k+1} = [x_k + \delta x_k]^{[0,y_{max}]} \tag{18}$$

$$\hat{\lambda}_{k+1} = [\hat{\lambda}_k + \alpha(\sum_{u\in\mathcal{U}} a_k^{(u)} - \sum_{t\in\mathcal{T}_k} p_{k-\tau_t})]^+ \tag{19}$$

where a new trace is started in slot k (added to $\mathcal{T}_k$) when x_k increases or when a trace completes and the number of active traces $|\mathcal{T}_k|$ falls below x_k.

Observe that $|\lambda_{g+1} - \lambda_g| \leq \alpha|\sum_{u\in\mathcal{U}} a_k^{(u)} - \sum_{t\in\mathcal{T}_k} p_{k-\tau_t}|$. Assuming that the number of packet arrivals $\sum_{u\in\mathcal{U}} a_k^{(u)}$ in slot k is bounded and step size α is sufficiently small that $|\lambda_{g+1} - \lambda_g| \leq \epsilon/(2Py_{max})$ then Lemma 1 can still be applied (the proof carries over essentially unchanged when using slots k rather than groups of slots g) to establish that $\hat{\lambda}_k$ is bounded provided y_{max} is sufficiently large. Similarly to the synchronised case it follows from (19) that,

$$\frac{1}{k}\sum_{i=1}^{k}(\sum_{u\in\mathcal{U}} a_i^{(u)} - \sum_{t\in\mathcal{T}_i} p_{i-\tau_t}) = \frac{1}{k}\sum_{i=1}^{k}\sum_{u\in\mathcal{U}} a_i^{(u)} - p_k^{\diamond} \leq \frac{\tilde{\lambda}_k}{\alpha k} \tag{20}$$

where $p_k^{\diamond} := \frac{1}{k}\sum_{i=1}^{k}\sum_{t\in\mathcal{T}_i} p_{i-\tau_t}$. Now, $\frac{1}{k}\sum_{i=1}^{k}\sum_{u\in\mathcal{U}} a_i^{(u)} \to \sum_{u\in\mathcal{U}} \bar{b}^{(u)}$ as $k\to\infty$. Since $\tilde{\lambda}_k$ is bounded then $\sum_{u\in\mathcal{U}} \bar{b}^{(u)} \leq p_k^{\diamond}$ as $k\to\infty$. That is, the mean service rate provided by the traces is sufficient to serve the mean rate of user packet arrivals and the input queue to the scheduler is stable.

V. Experimental Results

We built a prototype VPN implementation of the scheduler in (11)-(13) and also of its unsynchronised variant. The entire project including codes, scripts and datasets for all measurements in this paper is available at [7]. Using this VPN prototype, in this section we present experimental measurements of the scheduler performance under a wide range of conditions. In particular, while the analysis in the previous section focusses on throughput efficiency, our measurements allow us to also measure the delay performance of the scheduler and the impact of design choices (such as the use of unsynchronised trace activation) on this. We present results for both UDP and TCP traffic, a potential concern with TCP being possible interactions between the action of its congestion control and of the VPN scheduler (since both affect packet rate). In addition we present measurements for web page fetches carried out over the VPN and evaluate the delay and throughput efficiency achieved.

A. Hardware Setup

We begin by describing the hardware setup used. The network topology is shown schematically in Figure 8. Clients are connected to the VPN gateway A using an IPSec protected link. The link between nodes A and B (labelled Seculink) is the encrypted privacy-enhanced VPN tunnel. This is a public link, *i.e.* it is exposed to sniffing from adversaries. Seculink

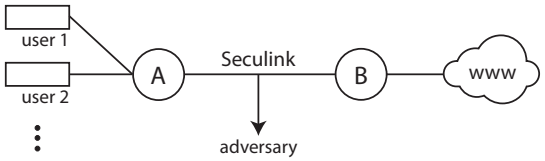

Fig. 8: Schematic illustrating the tunnel setup considered.

gateway A is a commodity server with an Intel(R) Core 2 Quad @2.66GHz CPU and 4GB RAM running Ubuntu 14.04.4 LTS. The other end of the tunnel, node B, is another commodity server with an Intel(R) Core 2 Duo @ 3.00GHz and 2GB RAM running Ubuntu 12.04.5 LTS. Nodes A and B use a mix of TP-LINK TG-3468 1000Mbps external ethernet cards and on-board gigabit ethernet cards (RTL8111 PCI Express and Intel 82567LM-3 respectively) for network connectivity. Traffic from node B is routed to the Internet via a campus gateway using a 100Mbps link (it is this link which limits the network rate since the tunnel between A and B is a gigabit connection). A NETGEAR JGS516 Gigabit switch is used to connect client machines to Seculink gateway A. The Seculink tunnel uses IPSec and traffic shaping to protect against DPI and traffic analysis attacks. Outgoing traffic from clients is sent to node B via node A and then forwarded by B to the campus gateway. Incoming traffic arriving at node B is forwarded to node A and then sent to the corresponding client.

B. Software Setup

Both Seculink and the links between clients and node A are protected with IPSec protocols AES-256 and SHA-256. Traffic shaping on Seculink is implemented using the `netfilter` and `netfilter_queue` libraries in C in conjunction with the `iptables` module in Linux. Gateway A queues packets from clients and transmits them over Seculink using our traffic shaping scheduler with slots of 1ms duration. When no client packets are queued for transmission yet the active traces require a packet to be sent then a dummy packet is generated and sent. Traffic received at node B is first filtered to remove dummy packets and remaining packets are then forwarded to the campus gateway over an unencrypted link. Responses received from the Internet are treated similarly, transmitted by node B to node A using our traffic shaping scheduler. In the UDP experiments traffic is generated using the PERIODIC method of the MGEN traffic generator. TCP traffic is generated by using `wget` to fetch dummy files of various sizes from a server located in the campus network.

C. Synchronised vs Unsynchronised Schedulers

With a synchronised scheduler traces only start and finish at the beginning of each cycle, a cycle being $9s$ duration in all of our tests unless otherwise stated. This means that when the arrival rate is changed in the middle of a cycle we have to wait until the start of the next cycle to update the number of active traces and so when new flows start they may experience significant delay. This is particularly an issue when no traces are active in the tunnel and a new client flow starts since in this case the scheduler cannot transmit any packets until the next

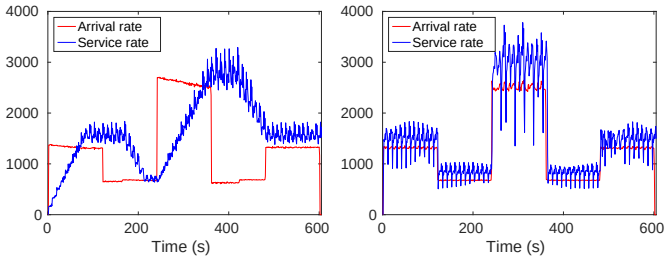

(a) Synchronised (b) Unsynchronised

Fig. 9: Comparison of tunnel rate adaptation in synchronised and unsynchronised scenarios.α in synchronised scenario. (duration: 10min, $\alpha = 1$, $P = 1682, n = 9615, c = 20, \gamma = 100$)

cycle starts (when some traces are already active the scheduler has the option to transmit some packets from the new flow e.g. by substituting for packets from other flows or for dummy packets). Similarly, when the rate of incoming traffic is less than the rate of a single trace then sufficient packets need to be queued before a trace is activated, causing delay.

To mitigate such delays on lightly loaded links, two unsynchronised modifications are implemented to "wake" the channel from silence upon sensing new incoming traffic:

1) *DNS Triggering*. Upon receiving a DNS packet when no traces are active, a new trace is immediately activated.
2) *Use of Running Average When Lightly-Loaded*. We maintain a running average of packet arrivals, $\bar{a} = (1-\zeta)\bar{a} + \zeta \sum_{u\in\mathcal{U}} a_k^{(u)}$, where $\zeta > 0$ is a design parameter. When $\bar{a}$ exceeds threshold a^* and no traces are active then a new trace is activated. In our experiments we use $\zeta = 0.001$, $a^* = 0.005$.

In addition, to avoid adding new traces in response to small spikes in arrivals, when $\gamma - \hat{\lambda}_k P < 0$ we only activate a new trace when also $\hat{\lambda}_k - \hat{\lambda}_{k-m} > 0$, where m is a parameter. In this way we only respond to a sustained increase in $\hat{\lambda}$. Similarly, when $\gamma - \hat{\lambda}_k P > 0$ we only decrease x_k when also $\hat{\lambda}_k - \hat{\lambda}_{k-m} < 0$. In our experiments we use $m = 100$ unless otherwise stated, which is found to provide a good compromise between responsiveness and sensitivity to small fluctuations (since slots are 1ms duration, this correponds to a window of 100ms duration).

To reduce sensitivity to small fluctuations in arrivals we also modify the $\tilde{\lambda}$ update to use $\max\{\bar{a}, \sum_{u\in\mathcal{U}} a_k^{(u)}\}$, namely:

$$\hat{\lambda}_{k+1} = [\hat{\lambda}_k + \alpha(\max\{\bar{a}, \sum_{u\in\mathcal{U}} a_k^{(u)}\} - \sum_{t\in\mathcal{T}_k} p_{k-\tau_t})]^+$$

In this way when the number of packet arrivals falls temporarily $\bar{a}$ is used instead of a_k.

Figure 9 illustrates the impact of these changes, comparing time histories of the arrival and services rates when using a synchronised scheduler vs an unsynchronised scheduler. It can be seen that the unsynchronised scheduler is much more responsive to changes to traffic load. Not only is the delay in increasing the service rate in response to increases in arrival rate reduced (so reducing the delay experienced by user packets) but also the delay in reducing the service rate

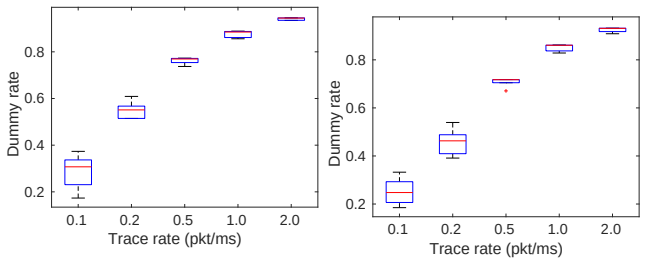

(a) Uplink (b) Downlink

Fig. 10: Dummy overhead vs trace rate for a sample website. For each rate data is shown for traces of duration 1, 2, 4, 8 and 16 seconds. ($\gamma = 4096$).

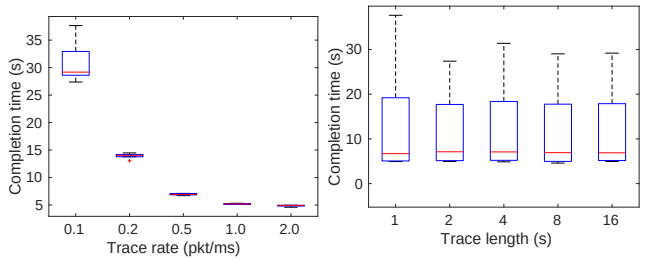

(a) Different rates (b) Different durations

Fig. 11: Illustrating the effect of trace rate and duration on the completion time of a sample web page. For each rate the data shown in (a) is for traces of duration 1, 2, 4, 8 and 16 second. Similarly, for each duration the data shown in (b) is for rates of 0.1, 0.2, 0.5, 1.0 and 2.0 pkts/ms. ($\gamma = 4096$).

in response to a fall in the arrival rate is also reduced (so reducing the number of dummy packets transmitted). Since it has better performance, in the rest of this section we will make use of this unsynchronised scheduler unless otherwise stated.

D. Choice of Trace Parameters

We make use of traces which transmit at a constant rate for a defined duration. Figure 10 shows measurements of the delay and fraction of dummy packets vs the trace rate and duration. This data is for an example web fetch using TCP and shows measurements for both the downlink and

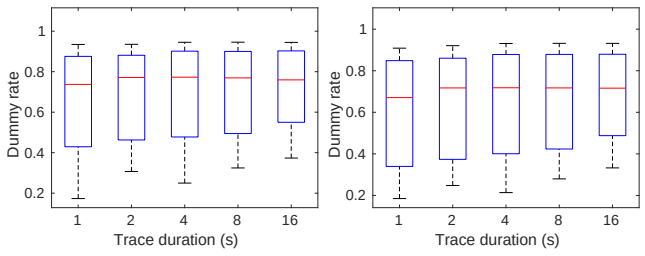

(a) Uplink (b) Downlink

Fig. 12: Illustrating the effect of trace duration on dummy overhead for a sample website. For each duration the data shown for rates of 0.1, 0.2, 0.5, 1.0 and 2.0 pkts/ms. ($\gamma = 4096$).

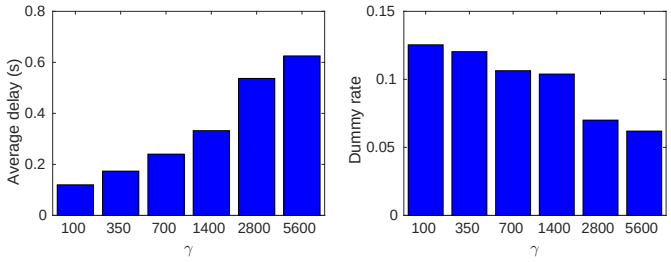

Fig. 13: Comparison between dummy and delay rates for CBR UDP traffic and different choices of γ. Given the negligible variance, average over 5 tests are presented for each γ. (duration: 5min, $\bar{a} = 2800$, $\alpha = 1$, $P = 1682, n = 9615, c = 20$). Values for each γ are averaged over 5 tests (error bars are not shown since they are too small to be clearly visible).

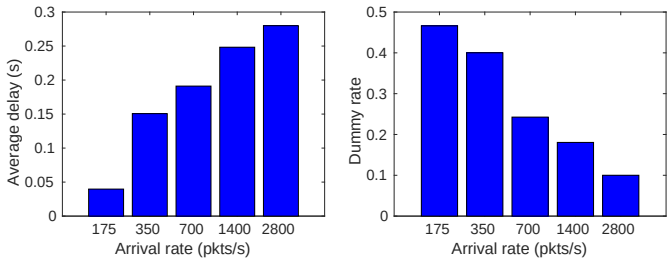

Fig. 14: Comparison between delay and dummy rates vs arrival rate for CBR UDP traffic with $\gamma = 1024$. (duration: 5min, $\alpha = 1$, $P = 1682, n = 9615, c = 20$). Values are averaged over 5 tests.

uplink (the uplink carries the TCP ACKs). Figure 11 shows the corresponding impact of the trace rate and duration on the web fetch completion time. It can be seen from these plots that increasing the trace rate reduces completion time but increases the dummy packet overhead, but the duration of trace has limited effect on completion time. The dummy packet overhead for different trace durations is also shown in Figure 12. From now on we use a trace of duration 9s and rate 0.2 pkts/ms, with aim of achieving a reasonable balance between delay and dummy packet overhead.

E. Performance with CBR UDP Traffic

In this section we study throughput efficiency and delay performance with constant rate UDP arrivals. By throughput efficiency we mean the mismatch between the transmit rate of the scheduler and the arrival rate of client packets. Recall that when there are insufficient client packets buffered at the scheduler then the scheduler transmits dummy packets so that its transmissions can continue to follow the predefined trace pattern. These dummy packets provide resistance to traffic analysis attacks but increase the load on the network and so we would like to minimise these. There is a fairly direct trade-off between delay and the volume of dummy packets transmitted since by increasing the backlog of client packets buffered at the input to the scheduler we reduce the need for dummy packets but increase the delay experience by client packets. Conversely, reducing the number of client packets buffered tends to increase the need for dummy packets but decrease delay.

This trade-off between throughput efficiency and delay can be seen in Figure 13, which plots measurements of the delay and dummy rate vs scheduler parameter γ which directly influences the queue backlog (with increases in γ increasing the backlog). The dummy rate value shown is the ratio of the dummy packets sent to the total number of packets sent (dummy plus processed packets). Note that there is no packet loss within the scheduler in these tests. It can be seen from Figure 13 that as γ is increased the delay rises but the rate at which dummy packets are sent falls. Observe, however, that even for relatively small values of γ the dummy rate remains reasonably small, which is encouraging.

Figure 14 shows corresponding measurements of delay and dummy rate as the arrival rate is varied (scheduler parameter γ is held constant at 1024 in these plots). It can be seen that the delay tends to increase with arrival rate and the dummy rate to fall. This can be understood by noting that as the arrival rate rises more user packets tend to be buffered at the scheduler. Hence, the delay rises but also user packets are more likely to be available when a transmisison opportunity occurs and so the number of dummy packets needed falls. Importantly, observe that the delay is consistently reasonable, rising to no more than 275ms at higher arrival rates, despite the extensive traffic shaping being carried out. Also, the decrease in dummy rate with increasing arrival rate means that the throughput efficiency increases with increasing load, so mitigating the dummy packet cost incurred by the traffic shaping.

F. Performance with TCP Flows

We now present data on the scheduler performance with TCP traffic (for the default Linux Cubic TCP variant). Figure 15 plots measurements of delay and fraction of dummy packets vs design parameter γ when fetching a 1024MB file from a campus server using `wget`. Data is shown for both the downlink and uplink, the packet rate on the uplink being roughly half of that on the downlink due to delayed acking by TCP. It can be seen that the delay increases with γ, as expected, but remains less than 100ms even for relatively large values of γ. The fraction of dummy packets is relatively insensitive to γ when using TCP.

Figure 16 plots the measured completion time vs file size fetched. For comparison, the corresponding data is also shown for a link with no traffic shaping. It can be seen that the cost, in terms of increase in completion time, is modest. Figure 17 provides additional detail, showing measured data on delay and dummy rate vs file size. It can be seen that the delay and dummy rate both tend to fall as the file size increases. The effect here is due to the time that it takes the scheduler to adapt to the arrival of a new flow: for longer flows this adaptation overhead gets washed out and amortised over many packets but for short flows its effect is more pronounced. This can be seen in Figure 18, which shows time histories of the number of active traces on both the downlink and uplink (uplink traces are carrying the TCP acks and so are fewer in number). The experiment is conducted by fetching a file of size 8192MB. The link reaches steady state in about 20 seconds.

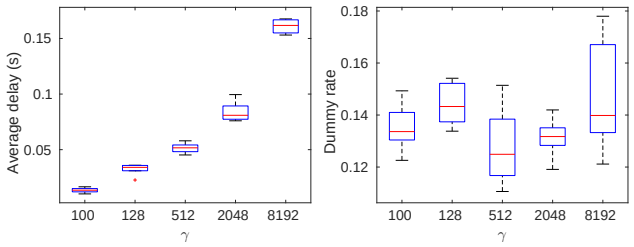

(a) Uplink

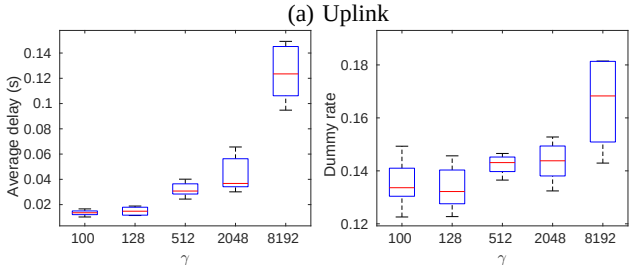

(b) Downlink

Fig. 15: Measurements of delay and fraction of dummy packets vs γ for TCP traffic. The data is measured fetching a 1024MB file, the average over 5 tests is presented for each γ ($\alpha = 1$, $P = 1682, n = 9615, c = 20$).

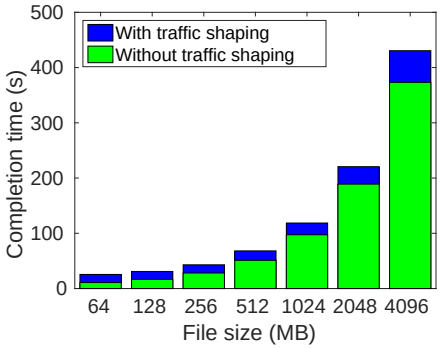

Fig. 16: Completion time vs file sizes when using TCP, $\gamma = 1024$. Results are shown both for the traffic shaped tunnel and for an unprotected link (no encryption, no traffic shaping). Values are averaged over 5 runs.

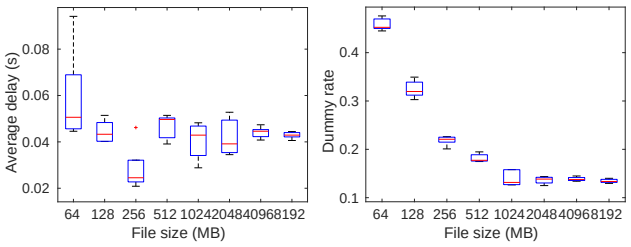

(a) Uplink

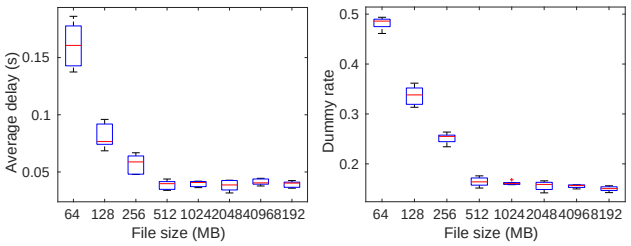

(b) Downlink

Fig. 17: Comparison between dummy and delay rate when using TCP to fetch files of different sizes, for $\gamma = 1024$. Box plots for values of 5 runs are shown.

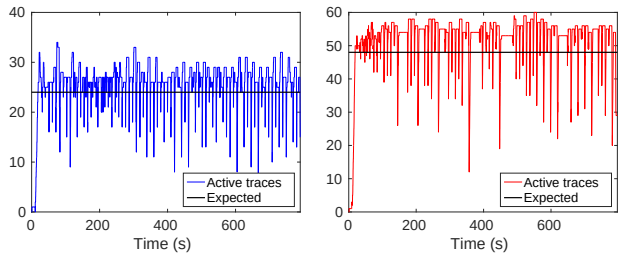

(a) Uplink (b) Downlink

Fig. 18: Illustrating convergence speed of active traces for uplink and downlink. The active traces approach channel capacity only 20 seconds into the fetch. Note that expected value of uplink capacity is half the value of downlink and that is due to TCP protocol sending one ACK for every two data packets.

G. Web Traffic Privacy

We fetched the home pages from the Alexa top 100 finance and health web sites in Ireland. Figure 19 shows four example trace time histories recorded during these fetches. It can be seen that the traces time histories in Figures 19a and 19c are identical and so evidently these two web pages cannot be distinguished by an attacker. Figure 20 plots the number of distinct trace time histories measured on the uplink and downlink while fetching the 100 web pages and also the number of pages for which each trace time history is observed. It can be seen that certain trace time histories are generated by 10-30 different web pages and so these web pages are indistinguishable to an attacker.

It can also be seen in Figure 20 that around 20 trace time histories are generated by a single web page, and so potentially vulnerable to attack. An example is shown in Figure 19d. Evidently the trace time history in Figure 19d cannot be distinguished from *two* fetches of the web sites in Figures 19a and 19c. The trace time history in Figure 19b is more complex, and raises the question of whether there exists one or more combinations of web page fetches that yield the same trace time history and so cannot be distinguished from this web fetch.

As an illustrative example we collected collected 10 samples for each of the two following browsing behaviours (i) a web page from our data set is fetched on a single tab and (ii) on three concurrent tabs, the words ”invisible”, ”history” and ”browsing” are queried from Google search with a 0.2 second gap between queries. The measured trace time histories are plotted in Figure 21. It can be seen that the time histories look very similar, thus allowing a user to plausibly deny that they fetched the single sensitive web page since the observed trace time history might equally have been generated by non-

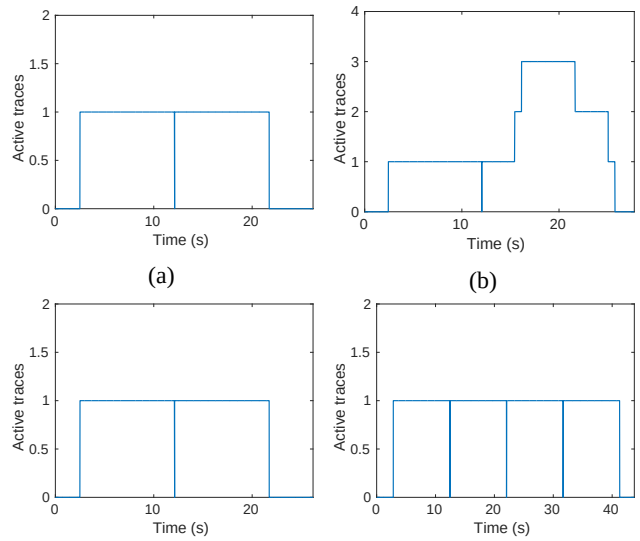

Fig. 19: Examples of trace time histories recorded when fetching the home pages of four websites from the Alexa top 100 finance and health web sites in Ireland. ($P = 1682, n = 9615, c = 20, \gamma = 1024$)

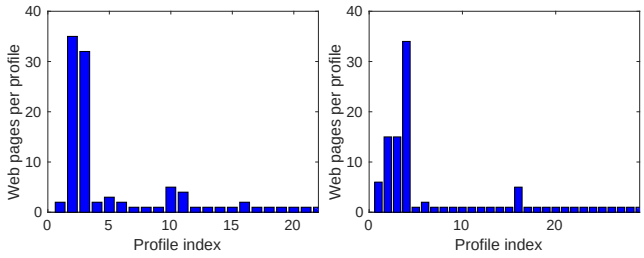

(a) Uplink (b) Downlink

Fig. 20: Counts of the number of web pages generating each distinct web trace time history.

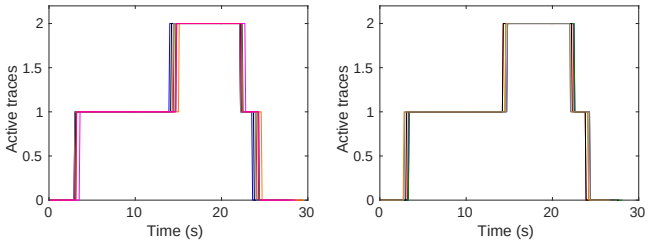

(a) Single web page (b) Concurrent web search

Fig. 21: Comparison between trace time histories from two different browsing sessions. In (a) a single web page is fetched while in (b) three Google search queries fare conducted on concurrent tabs with a 0.2s time gap between fetches. ($P = 1682$, $n = 9615$, $c = 20$, $\gamma = 1024$).

sensitive google queries.

To get more insight, we proceed as follows. Let $\mathcal{H}$ denote the set of measured trace time histories. For trace time history $h \in \mathcal{H}$ we determine (by exhaustive search) the combinations $C_{h,1}, C_{h,2}, \ldots, C_{h,m_h}$ of the other trace

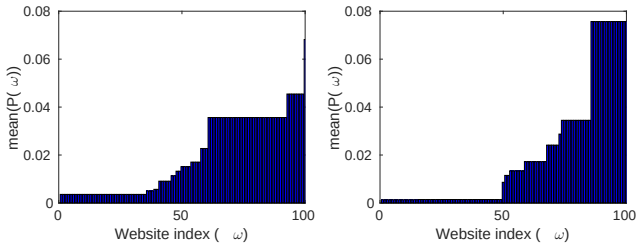

(a) Uplink (b) Downlink

Fig. 22: Probability $\mathrm{P}(\omega)$ that web page ω was fetched given an observed trace time history vs ω and assuming equiprobable combinations.

time histories that are indistinguishable from h. Combination $C_{h,i} = \{(h_{i,1}, n_1), (h_{i,2}, n_2), \ldots\}$ with $h_{i,j} \in \mathcal{H} \setminus \{h\}$ and n_j the number of times that $h_{i,j}$ is repeated in the combination. Let $\mathcal{W}_h$ denote the set of web pages that generate trace time history h, so $|\mathcal{W}_h|$ is the number of web pages that generate h (see Figure 20). Let the web page fetched W be a random variable that takes values in $\cup_{h \in \mathcal{H}} \mathcal{W}_h$. Assume, for simplicity, that the web pages in $\mathcal{W}_h$ are equally likely to be fetched. Then when combination $C_{h,i}$ is observed the probability that page ω was fetched is,

$$\mathrm{P}(W = \omega | C_{h,i}) = \sum_{(h,n) \in C_i : \omega \in \mathcal{W}_h} \sum_{j=1}^{n} \binom{n}{j} \frac{1}{|\mathcal{W}_h|}^{j} (1 - \frac{1}{|\mathcal{W}_h|})^{n-j} \quad (21)$$

Let C be a random variable which is the combination used. Assume, again for simplicity, that each combination $C_{h,1}, C_{h,2}, \ldots, C_{h,m_h}$ is equally probable when trace h is observed, i.e $\mathrm{P}(C = C_{h,i}) = 1/m_h$. Then the probability that web page ω was fetched given that trace time history h was observed is,

$$\mathrm{P}(W = \omega | h) = \sum_{i=1}^{m_h} \mathrm{P}(W = \omega | C_{h,i}) / m_h \quad (22)$$

When the traces $h \in \mathcal{H}$ are equally likely to be observed then $\mathrm{P}(W = \omega) = \frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} \mathrm{P}(W = \omega | h)$. Figure 22 plots $\mathrm{P}(W = \omega)$ calculated using (22) for each of the 100 web pages, sorted in increasing order. It can be seen that this probability is less than 0.08 on the downlink and less than 0.05 on the uplink. We can conclude therefore that with the foregoing assumptions the user can reasonably deny that page ω was fetched despite an attacker observing the transmitted packet trace.

To get a sense of the sensitivity of these values to the assumption of equiprobability, for comparison, let $h^* \in \arg\min_{h \in \mathcal{H}} \mathrm{P}(W = \omega | h)$ and suppose that this worst case trace h^* is always observed i.e. $\mathrm{P}(W = \omega) = \mathrm{P}(W = \omega | h^*)$. Figure 23 shows the calculated $\mathrm{P}(W = \omega | h^*)$ for our measured web fetches. It can be seen that $\mathrm{P}(W = \omega | h^*)$ is higher than in Figure 22, as expected, and indeed reaches one but only for a small number of web pages. Hence, even in this worst case a user has fairly level of strong deniability.

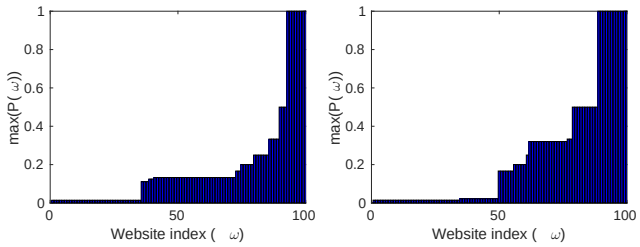

(a) Uplink (b) Downlink

Fig. 23: Probability P(ω) that web page ω was fetched given an observed trace time history vs ω and assuming worst combination has probability one.

VI. Conclusions

In this paper we introduce a trace-based tunnel that is resistant to traffic analysis in the sense that it provides deniability to users that any specified web page was fetched given that a specified packet trace is observed on the tunnel. We present a scheduler design for managing the transmission of traces to satisfy user traffic demand while maintaining reasonably low delay and throughput overhead due to dummy packets. Experimental results are also presented demonstrating the effectiveness of this scheduler under a range of realistic network conditions and real web page fetches.

Appendix: Proof of Lemma 1

Proof: We begin by observing that,

$$\begin{aligned}
&\arg\min_{x\in[-1,1]}(\gamma-\lambda_g P)x \\
&=\arg\min_{x\in[-1,1]}(\gamma-\lambda_g P)(\bar{y}_g+\beta x)+\lambda_g\sum_{u\in\mathcal{U}}\bar{b}^{(u)} &(23)\\
&=\arg\min_{x\in[-1,1]}L(\bar{y}_g+\beta x,\lambda_g) &(24)
\end{aligned}$$

Update (7) projects $\bar{y}_g+\beta x$ onto interval $[0,y_{max}]$. Combining this projection, therefore, into the optimisation we have that steps (6)-(7) can be rewritten equivalently as

$$\bar{y}_{g+1}=\bar{y}_g+\beta\arg\min_{\substack{x\in[-1,1]:\\ \bar{y}_g+\beta x\in[0,y_{max}]}}L(\bar{y}_g+\beta x,\lambda_g) \qquad (25)$$

We proceed by considering two cases.

Case (i):

$$|L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)|\geq\epsilon \qquad (26)$$

where $x\in\arg\min_{\substack{x\in[-1,1]:\\ \bar{y}_g+\beta x\in[0,y_{max}]}}L(\bar{y}_g+\beta x,\lambda_g)$. Note that since $(\gamma-\lambda_g P)\beta x\leq 0$ this case corresponds to $L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)\leq-\epsilon$. Then,

$$\begin{aligned}
&L(\bar{y}_{g+1},\lambda_{g+1})-L(\bar{y}_g,\lambda_g)\\
&=L(\bar{y}_g+\beta x,\lambda_{g+1})-L(\bar{y}_g,\lambda_g) &(27)\\
&=(\gamma-\lambda_{g+1}P)(\bar{y}_g+\beta x)-(\gamma-\lambda_g P)\bar{y}_g\\
&\qquad-(\lambda_g-\lambda_{g+1})\sum_{u\in\mathcal{U}}\bar{b}^{(u)} &(28)\\
&=(\gamma-\lambda_g P)\beta x+(\lambda_g-\lambda_{g+1})(P(\bar{y}_g+\beta x)-\sum_{u\in\mathcal{U}}\bar{b}^{(u)})\\
&\leq-\epsilon+(\lambda_g-\lambda_{g+1})(P(\bar{y}_g+\beta x)-\sum_{u\in\mathcal{U}}\bar{b}^{(u)}) &(29)
\end{aligned}$$

Now, $|\lambda_{g+1}-\lambda_g|\leq\epsilon/(2(Py_{max}-b^{\diamond}))$ and so

$$L(\bar{y}_{g+1},\lambda_{g+1})-L(\bar{y}_g,\lambda_g)\leq-\epsilon+\epsilon/2\leq-\epsilon/2 \qquad (30)$$

That is, the Lagrangrian is strictly decreasing.

Case (ii):

$$|L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)|<\epsilon \qquad (31)$$

Since $(\gamma-\lambda_g P)\beta x\leq 0$ this case corresponds to $-\epsilon<L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)\leq 0$. Hence, $L(\bar{y}_{g+1},\lambda_{g+1})-L(\bar{y}_g,\lambda_g)=(\gamma-\lambda_g P)\beta x+(\lambda_g-\lambda_{g+1})(P(\bar{y}_g+\beta x)-\sum_{u\in\mathcal{U}}\bar{b}^{(u)})\leq\epsilon/2$.

Thus,

$$\begin{aligned}
&L(\bar{y}_{g+1}+\beta x,\lambda_{g+1})-L(\bar{y}_{g+1},\lambda_{g+1})\\
&\geq L(\bar{y}_{g+1}+\beta x,\lambda_{g+1})-L(\bar{y}_g,\lambda_g)-\epsilon/2 &(32)\\
&=(\gamma-\lambda_{g+1}P)(\bar{y}_{g+1}+\beta x)-(\gamma-\lambda_g P)\bar{y}_g-\epsilon/2 &(33)\\
&=(\gamma-\lambda_g P)(\bar{y}_{g+1}-\bar{y}_g)+(\gamma-\lambda_g P)\beta x-\epsilon/2\\
&\quad+(\lambda_g-\lambda_{g+1})P(\bar{y}_{g+1}+\beta x) &(34)\\
&\geq-\epsilon\beta-\epsilon\beta-\epsilon/2-\epsilon/2=-2\epsilon\beta-\epsilon=-(2\beta+1)\epsilon &(35)
\end{aligned}$$

$L(\bar{y}_g,\lambda_g)$ is strictly decreasing when $|L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)|\geq\epsilon$ and once less than ϵ this difference increases to at most $(2\beta+1)\epsilon$. Hence, there exists a finite $\bar{g}$ such that for $g=\bar{g}$ then $-(2\beta+1)\epsilon<L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)\leq 0$ and once this happens this inequality holds for all larger values of $g>\bar{g}$. Now $L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)=(\gamma-\lambda_g P)\beta x$ and so $-(2\beta+1)\epsilon<(\gamma-\lambda_g P)\beta x\leq 0$. Provided $x\in\{-1,1\}$ when $|L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)|<(2\beta+1)\epsilon$ then it follows that $|\gamma-\lambda_g P|\leq(2\beta+1)\epsilon/\beta$ for all $g\geq\bar{g}$ as claimed.

It remains to show that $x\in\{-1,1\}$ when $|L(\bar{y}_g+\beta x,\lambda_g)-L(\bar{y}_g,\lambda_g)|<(2\beta+1)\epsilon$ and y_{max} is sufficiently large. To see this, suppose that $0<x<1$ (it is enough to consider positive valued updates x since our interest is in the right-hand limit of interval $[0,y_{max}]$). Increasing y_{max} to $y_{max}+\beta$ leads to $x=1$ since the optimisation in (25) is linear and so the solution lies at an extreme point. At the next iteration k we may need to increase y_{max} again, but this process necessarily terminates in finite time because λ_g increases by $\alpha(\sum_u\bar{b}^{(u)}-\bar{y}_g P)$ at each iteration and so eventually either (a) force $\gamma-\lambda_g P$ to change sign and so x also change sign, or (b) λ_g will start to decrease (since $\bar{y}_g$ increases and $\bar{b}^{(u)}$ is bounded, it is here that $y_{max}>b^{\diamond}$ is needed) and so $x=1$ remains feasible. ■

REFERENCES

[1] G. D. Bissias, M. Liberatore, D. Jensen, and B. N. Levine. Privacy Vulnerabilities in Encrypted HTTP Streams. In G. Danezis and D. Martin, editors, *Privacy Enhancing Technologies*, volume 3856 of *Lecture Notes in Computer Science*, pages 1–11. Springer Berlin Heidelberg, 2006.

[2] X. Cai, X. Ch. Zhang, B. Joshi, and R. Johnson. Touching from a Distance: Website Fingerprinting Attacks and Defenses. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, pages 605–616, New York, NY, USA, 2012. ACM.

[3] Xiang Cai, Rishab Nithyanand, and Rob Johnson. Cs-buflo: A congestion sensitive website fingerprinting defense. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society*, WPES '14, pages 121–130, New York, NY, USA, 2014. ACM.

[4] Xiang Cai, Rishab Nithyanand, Tao Wang, Rob Johnson, and Ian Goldberg. A systematic approach to developing and evaluating website fingerprinting defenses. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, pages 227–238, New York, NY, USA, 2014. ACM.

[5] S. Chen, R. Wang, X. Wang, and K. Zhang. Side-channel leaks in web applications: A reality today, a challenge tomorrow. In *2010 IEEE Symposium on Security and Privacy*, pages 191–206, May 2010.

[6] K. P. Dyer, S. E. Coull, T. Ristenpart, and T. Shrimpton. Peek-a-Boo, I Still See You: Why Efficient Traffic Analysis Countermeasures Fail. In *Security and Privacy (SP), 2012 IEEE Symposium on*, pages 332–346, May 2012.

[7] S. Feghhi. Seculink Tunnel Project: Codes and Measurements, 2016. available at: www.scss.tcd.ie/disciplines/stable/seculink.

[8] S. Feghhi and D. Leith. A Web Traffic Analysis Attack Using Only Timing Information. *IEEE Transactions on Information Forensics and Security*, 11(8):1747–1759, April 2016.

[9] Tobias Flach, Nandita Dukkipati, Andreas Terzis, Barath Raghavan, Neal Cardwell, Yuchung Cheng, Ankur Jain, Shuai Hao, Ethan Katz-Bassett, and Ramesh Govindan. Reducing web latency: the virtue of gentle aggression. *ACM SIGCOMM Computer Communication Review*, 43(4):159–170, 2013.

[10] Xun Gong, Negar Kiyavash, and Nikita Borisov. Fingerprinting websites using remote traffic analysis. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, CCS '10, pages 684–686, New York, NY, USA, 2010. ACM.

[11] Dominik Herrmann, Rolf Wendolsky, and Hannes Federrath. Website fingerprinting: Attacking popular privacy enhancing technologies with the multinomial naïve-bayes classifier. In *Proceedings of the 2009 ACM Workshop on Cloud Computing Security*, CCSW '09, pages 31–42, New York, NY, USA, 2009. ACM.

[12] A. Hintz. Fingerprinting Websites Using Traffic Analysis. In Roger Dingledine and Paul Syverson, editors, *Privacy Enhancing Technologies*, volume 2482 of *Lecture Notes in Computer Science*, pages 171–178. Springer Berlin Heidelberg, 2003.

[13] M. Liberatore and B. N. Levine. Inferring the Source of Encrypted HTTP Connections. In *Proceedings of the 13th ACM Conference on Computer and Communications Security*, CCS '06, pages 255–263, New York, NY, USA, 2006. ACM.

[14] L. Lu, E. C. Chang, and M. C. Chan. Website Fingerprinting and Identification Using Ordered Feature Sequences. In D. Gritzalis, B. Preneel, and M. Theoharidou, editors, *Computer Security ESORICS 2010*, volume 6345 of *Lecture Notes in Computer Science*, pages 199–214. Springer Berlin Heidelberg, 2010.

[15] X. Luo, P. Zhou, E. W. W. Chan, W. Lee, R. K. C. Chang, and R. Perdisci. HTTPOS: Sealing information leaks with browser-side obfuscation of encrypted flows. In *In Proc. Network and Distributed Systems Symposium (NDSS). The Internet Society*, 2011.

[16] B. Miller, L. Huang, A. D. Joseph, and J. D. Tygar. I Know Why You Went to the Clinic: Risks and Realization of HTTPS Traffic Analysis. In Emiliano De Cristofaro and Steven J. Murdoch, editors, *Privacy Enhancing Technologies*, volume 8555 of *Lecture Notes in Computer Science*, pages 143–163. Springer International Publishing, 2014.

[17] Rishab Nithyanand, Xiang Cai, and Rob Johnson. Glove: A bespoke website fingerprinting defense. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society*, WPES '14, pages 131–134, New York, NY, USA, 2014. ACM.

[18] A. Panchenko, L. Niessen, A. Zinnen, and T. Engel. Website Fingerprinting in Onion Routing Based Anonymization Networks. In *Proceedings of the 10th Annual ACM Workshop on Privacy in the Electronic Society*, WPES '11, pages 103–114, New York, NY, USA, 2011. ACM.

[19] S.Souders. Velocity and the bottom line. In *Velocity (Web Performance and Operations Conference)*, 2009.

[20] Q. Sun, D. R. Simon, Yi-Min Wang, W. Russell, V. N. Padmanabhan, and Lili Qiu. Statistical identification of encrypted web browsing traffic. In *Security and Privacy, 2002. Proceedings. 2002 IEEE Symposium on*, pages 19–30, 2002.

[21] W. Tao, C. Xiang, N. Rishab, R. Johnson, and I. Goldberg. Effective Attacks and Provable Defenses for Website Fingerprinting. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 143–157, San Diego, CA, August 2014. USENIX Association.

[22] V. Valls and D.J. Leith. Max-weight revisited: Sequences of nonconvex optimizations solving convex optimizations. *Networking, IEEE/ACM Transactions on*, PP(99):1–1, 2015.

[23] Tao Wang and Ian Goldberg. Walkie-talkie: An effective and efficient defense against website fingerprinting. Technical report, Technical Report 2015-09, CACR, 2015. http://cacr. uwaterloo. ca/techreports/2015/cacr2015-09. pdf, 2015.

[24] C. V. Wright, S. E. Coull, and F. Monrose. Traffic morphing: An efficient defense against statistical traffic analysis. In *In Proceedings of the 16th Network and Distributed Security Symposium*, pages 237–250. IEEE, 2009.

[25] Yang Xue and E. Y. Vasserman. Simple and compact flow fingerprinting robust to transit through low-latency anonymous networks. In *2016 13th IEEE Annual Consumer Communications Networking Conference (CCNC)*, pages 765–773, Jan 2016.

[26] Wenxuan Zhou, Qingxi Li, Matthew Caesar, and P Godfrey. ASAP: A low-latency transport layer. In *Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies*, page 20. ACM, 2011.